

GOT RISK? A RISK-CENTRIC PERSPECTIVE FOR SPACECRAFT TECHNOLOGY DECISION-MAKING

Martin S. Feather, Steven L. Cornford & Kelly Moran
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA 91109-8099
{Martin.S.Feather, Steven.L.Cornford, Kelly.Moran}@Jpl.Nasa.Gov

ABSTRACT

A risk-based decision-making methodology conceived and developed at JPL and NASA has been used to aid in decision making for spacecraft technology assessment, adoption, development and operation. It takes a risk-centric perspective, through which risks are used as a reasoning step to interpose between mission objectives and risk mitigation measures. The novel aspects of this methodology lie in:

Broad-ranging treatment of objectives, risks and mitigations: objectives encompass science objectives of the mission, development considerations, and constraints on operation; risks are broadly defined to include the risk of failing to design a system with adequate performance, compatibility and robustness in addition to more traditional implementation and operational risks; mitigations include architectural and design choices, technology plans and technology back-up options, test-bed and simulation options, engineering models and hardware/software development techniques and other more traditional risk reduction techniques (tests, analyses, inspections, etc).

Quantitative treatment of the relationships among these concepts: risks are quantitatively related to objectives to indicate the extent to which each risk, were it to occur, would detract from attainment of that objective; mitigations are quantitatively related to risks to indicate the extent to which each mitigation, were it to be applied, would reduce the likelihood and/or impact of the risk (or, in some cases, *increase* the risk).

A software-supported process to gather, scrutinize and reason in terms of the objectives/risks/mitigations information: custom software facilitates all the steps of this

process. This makes feasible the consideration of hundreds of items and thousands of quantitative connections among them – essential to informed decision making in the technology-rich yet risk averse setting in which space missions are planned.

Applications to NASA spacecraft technologies have demonstrated: improved insights into a variety of risks, ability to trade and calibrate risk across discipline boundaries, optimized planning of how to address risk, risk-informed comparison among design alternatives, and risk-guided descopeing (strategic abandonment of objectives, necessitated when the risk-based analysis reveals the infeasibility of satisfactory levels of objective attainment with the resources available).

1 Introduction – Defect Detection and Prevention

At JPL and NASA we have been developing and applying a risk-based approach to assist early-lifecycle planning of complex system developments. The approach is called “Defect Detection and Prevention” (DDP), the name reflecting its origins as a method intended for quality assurance planning of hardware systems [Cornford 1998]. Various aspects of DDP have been described in previously published papers, e.g., overviews of its status and application are in [Cornford et al, 2001]; the look and feel of the tool support in [Feather et al, 2000]. More information can be obtained from the DDP website: <http://ddptool.jpl.nasa.gov>

In this paper we focus on the significant differences between DDP and more traditional forms of risk analysis. The origin of these differences is rooted in the intended purpose of DDP, namely as a decision-

making aid during planning and development of spacecraft technologies. We are especially focused on decision-making in the early stages of development. This is an important but challenging time of the life cycle. It is important because these early decisions have the most leverage to influence the development to follow. It is challenging because information on which to base those decisions is incomplete and uncertain, and in the case of advanced technologies and systems there is little past experience from which to extrapolate.

The remainder of the paper is organized as follows:

Section 2 presents an overview of DDP. Sections 3 - 6 consider how DDP differs from conventional practices, and what are the comparative benefits that accrue from the DDP approach. More specifically: Section 3 focuses on DDP's objectives, Section 4 focuses on DDP's risks, Section 5 focuses on DDP's mitigations, and Section 6 focuses on DDP's Human-Computer Interface. Section 7 briefly summarizes results of DDP applications. Section 8 concludes with some consideration of the drawbacks of DDP, and future work.

2 DDP Overview

2.1 DDP Objects

The DDP process deals with three key sets of objects: "*Objectives*", "*Risks*" and "*Mitigations*".

Objectives (a.k.a. Requirements) are the things that the system is to achieve, the limitations on how it is to be developed, and restrictions on how it must operate. Objectives are assigned different "weights" to reflect their relative importance.

Risks are all the kinds of things that, should they occur, would lead to failure to attain Objectives. Note that this use of the word "risk" may appear somewhat non-standard. The usual definition of risk, as "probability * severity" is a measure we calculate from the DDP information and

associate with these. Depending on the circumstances, we have alternately referred to these DDP objects as "risk elements", "failure modes", "defects" or "obstacles" to more intuitively reflect their nature. In this paper, we will refer to them as simply "risks".

Risks are assigned an "a-priori" likelihood, namely the likelihood of that Risk occurring if nothing is done to prevent it. Risks are also assigned a cost of "repair" (a.k.a. "correction"). This is the cost of repairing the problem – for example, the cost of repairing a component damaged during assembly, or the cost of correcting a software bug introduced during coding and detected at test time. It is often the case that these repair costs escalate through the course of the development lifecycle (e.g., the cost of correcting a design flaw if that flaw is detected at design time vs. during implementation vs. during system test vs. after release). In the DDP model Risks are assigned time-specific repair costs that can be used to capture this escalation.

Mitigations are the options that could be taken to prevent or reduce Risks. These could be training, adoption of standards, tests, analyses, inspections, reviews, redundant design elements, etc. Mitigations encompass preventative measures (which decrease the likelihood of risks arising in the first place), detections (which if applied before use enable problems to be identified, and repaired), and alleviations (which reduce the impact of problems should they occur). Note that our use of the term "mitigation" encompasses all three: we use the term "alleviation" to refer specifically to actions that decrease only the consequence of a risk.

Each Mitigation is assigned a cost, namely the resource costs of applying it. In our world of spacecraft development, there are typically several kinds of critical resources, e.g., budget (\$), mass, volume, electrical power. Each Mitigation is also assigned a time, typically the "phase" in the development effort at which it is applied (e.g., requirements time, design time, coding time). It is possible to use other time scales (e.g., financial quarters or, for long duration developments, years).

2.2 DDP Relationships

The DDP process deals with quantitative relationships that link Objectives, Risks and Mitigations, as follows:

Impacts are the quantitative relationships between Objectives and Risks, namely the proportion of the objective attainment that would be lost should the Risk occur. A Risk can impact multiple Objectives to different extents, and similarly an Objective can be impacted by multiple Risks, again to different extents.

Effects are the quantitative relationships between Mitigations and Risks, namely the proportion by which a Mitigation reduces a Risk should that Mitigation be applied. A Mitigation can effect multiple Risks to different extents, and similarly a Risk can be effected by multiple Mitigations, again to different extents.

2.3 DDP overall

The overall form of a DDP model is sketched in Fig. 1.

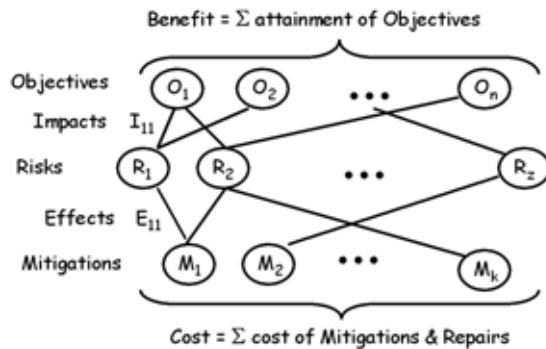


Figure 1. Topology of DDP model

In almost all cases the cost of all possible Mitigations exceeds the resources available. The primary purpose for which DDP was constructed is to help in the judicious selection of which Mitigations to perform, so as to minimize overall risk (and therefore maximize attainment of Objectives) while remaining within the constraints on the resources available. In practice there are other significant benefits that can be gained by application of DDP. We will point these out in the course of the paper.

The challenging nature of problems we face is evident in Fig. 2, which shows the topology of the DDP information for an actual study of a spacecraft technology. Custom software supports the application of DDP, enabling models of this complexity to be constructed and effectively utilized.

2.4 DDP example

As illustration, we summarize an actual application of DDP. In order to avoid proprietary issues, we do not reveal specifics.

The example is the application of DDP to an advanced technology for data storage, intended for spacecraft use. The technology had been demonstrated successfully in a laboratory setting. The purpose of the DDP study was to help plan the next step in the development of the technology towards actual spacecraft use. To do this, it was important to understand the science needs driving the demand for data storage, the mission environment in which the technology would have to operate, and the resource constraints




Figure 2. Topology of the data in a DDP application

on the technology implementation.

The **Objectives** fell into the following areas:

- Mission characteristics: e.g.,
 - how long would this data storage technology have to survive?
 - how much radiation must the technology tolerate?
 - what temperature range must the technology be able to survive?
 - what levels of vibration and shock must the technology be able to survive?
- Data characteristics: e.g.,
 - how much data would it be required to store?
 - how quickly must it be able to accept data (e.g., for a fly-by mission, the data input rate needed could be very high)?
 - how much tolerance is there for data error?
 - how many read-write cycles must it support?
- Resource characteristics: e.g.,
 - How much power is available to the unit?
 - How big can it be?

Over 30 detailed objectives were gathered. Similarly, **Risks** relevant to these objectives were gathered, encompassing risk areas of:

- Radiation induced problems
- Contamination induced problems
- Temperature induced problems
- Materials degradation
- Hazards prior to launch (e.g., Electrostatic Discharge; humidity)
- Technology packaging issues
- Technology-specific issues (we apologize for the necessity to remain silent on the details of these)

Almost 60 detailed risks were identified.

Finally, **Mitigation** options for preventing or reducing risks were gathered. These fell into areas of:

- Preventative measures: e.g.,

- Design choices (e.g., use of radiation shielding; vibration isolation)
- Development practices and procedures (e.g., store in dry Nitrogen)
- Analyses: e.g.,
 - Sensitivity analyses of the technology itself
 - Mission analyses of the environment in which the technology must operate
- Tests and prototypes: e.g.,
 - Build a “breadboard” prototype
 - The gamut of spacecraft tests (e.g., temperature cycling, vibration, radiation), for each of which there are typically a number of variants
 - Functional tests of the technology itself

Over 80 such items were identified. For purposes of comparison, the option of “flight validation” (i.e., build a working unit, and flying it on a spacecraft mission to test it, but not rely upon it) was also included. Such flight validation, while effective as a way of revealing the presence of spaceflight-relevant problems, is of course very expensive. Its inclusion served as a benchmark against which to compare combinations of alternative measures, to see how well they would reduce risks, and how much they would cost.

These objectives, risks and mitigation options were gathered in sessions in which there were present representatives of mission scientists, engineering experts, quality assurance personnel, operations, management, etc.

These items were listed, and then correlated with one another in terms of the “**impact**” and “**effect**” relationships described earlier. For each Risk item the team identified which of the Objectives the occurrence of that Risk would detract from, and by how much. Over 300 such impact relationships were identified (it is typical that a Risk impacts only some of the Objectives, hence the result of such information gathering is a relatively sparse

matrix of non-zero correlations between Risks and Objectives). These quantitative estimates of impacts were generally expressed to only one significant digit – as it turns out, such modest precision is quite sufficient to guide key decision-making, and indeed it is unlikely that more precise information is available when considering the novel application of an advanced technology. The same phenomenon is true of the quantitative estimates of the effectiveness of Mitigations at preventing/reducing Risks.

3 Objectives

In this section we consider DDP's explicit representation of multiple Objectives – how it differs from conventional practices, and what are the comparative benefits that accrue from the DDP approach.

3.1 Conventional Practice

In most conventional risk assessment and risk management methods, risk is assessed against a small number of criteria. For example, Probabilistic Risk Assessment (PRA) methods are able to accommodate several alternate “end states”, and this capability is used to represent, say, several major alternative outcomes, so as to be able to assess the likelihoods of each. Early lifecycle risk assessment methods are commonly used to assess risks against a small number of major concerns (e.g., cost – will the development be completed within budget?; schedule – will the development be completed in time?; function – will the system that results achieve its function?).

3.2 DDP's Treatment of Objectives

In DDP, Objectives are user-defined entities, allowing (indeed, encouraging) the representation of numerous such entities. In our application of DDP to the study of advanced technologies, we have commonly listed dozens of separate Objectives. The benefits of this are as follows:

Capture of a wide variety of concerns.

For example, in a study that focused on development of a novel memory technology, objectives included functional performance

objectives (e.g., how much data it could hold, how quickly it could read it), system concerns (e.g., how much power it would consume, how large it would be), environmental concerns (e.g., conditions under which it would operate), and an objective stated as “No other technology is better” – this last one was there to capture the important aspect of competitive advantage: why fund the development of a novel technology if, by the time it would be ready for use, there would likely be some other superior technology available?

Capture varying levels of detail commensurate with the problem at hand.

For example, in a study that focused on assessment of an electronics packaging technique's suitability to perform reliably while withstanding the cyclic temperature extremes of the Martian surface environment, one Objective said simply “Mass”, there to capture the importance of the mass of the package; in contrast, there were seven detailed Objectives concerning the issue of surviving the temperature cycles (the switch between high and low temperatures during every Martian day/night cycle). This disparity was a deliberate reflection of the issues the study needed to focus on. Mass was recognized to be a well-understood issue, so it sufficed to capture mass concerns within a single Objective, while temperature cycling was at the heart of the study, and so warranted detailed elaboration.

Trace Risks back to the specific Objectives those Risks threaten. This provides the ability to discern those Objectives proving the most problematic to attain (i.e., those Objectives threatened by Risks which in turn are proving to be expensive to reduce). For example, in one study, a particularly problematic performance Objective turned out to be a suggestion that the scientists on hand in the study quickly confirmed to be unnecessarily stringent. The less demanding Objective they were able to offer as its replacement was much easier to attain. Another utilization of this capability is to allow consideration of “descope” options –

by which we mean abandonment (or reprioritization) of Objectives. Such descoping is necessary when the available resources are insufficient to permit satisfactory attainment of the initial set of Objectives. Descoping may take the form of a less ambitious design that achieves less, but whose implementation is feasible.

Reprioritization – for example, swapping the primary and secondary mission goals – is a less drastic form of descoping in which problematic Objectives are retained, but their importance is decreased so that their attainment is no longer the driving factor. In our studies of advanced technologies descoping is useful when there are multiple options for how a technology might be applied. The flexibility it conveys allows us to use DDP studies to help locate the ideal application of technologies [Feather et al, 2002].

4 Risks

In this section we consider DDP's treatment of Risks – how it differs from conventional practices, and what are the comparative benefits that accrue from the DDP approach.

4.1 Conventional Practice

Most risk assessment processes take as starting point a design, and focus on the risks *remaining* in that design. The benefits of the steps that were (or are planned to be) followed during design and implementation are reflected in the reliabilities assigned to its components. For example, consider navigation software: if its programmers are skilled in development of such software, then the number coding defects due to domain misunderstandings present in the code prior to commencement of testing may reasonably be assumed to be small.

4.2 DDP's Treatment of Risks

In DDP, the steps that lead to each Risk's status in the design are explicitly represented. Typically a Risk begins with an "a-priori" likelihood of 1 (certainty), and it is only through the risk-reducing effects of the

planned design steps that its DDP-computed likelihood will become small. Exceptions to this are Risks that are rooted in natural phenomena (e.g., the risk of lightning strike during the launch window), which have a likelihood of less than 1. The benefits of DDP's treatment of Risks are as follows:

Capture the assumptions that underpin the risk assessment. For example, the assessment of certain Risks may assume that there will be sufficient time for running a very thorough suite of tests. In the event that schedule slips preclude such extensive testing, it will be useful to know which Risks' assessments were contingent upon those tests. In practice we vary the extent to which we record the details of these assumptions. At one extreme we simply make note of them as textual comments appended to Mitigations, so that on future occasions we can at least review those notes to check whether the assumptions are still valid. At the other extreme, we represent (for example) the testing assumptions as DDP "Mitigations", allowing us to reason about the effect of performing, or not performing, them (see section 5 for further details).

Support decision making to help choose which (out of potentially many) Mitigations to employ. This is only useful when the risk assessment is performed early in the development, while there is still time to choose. Recall that we are especially focused on decision-making in the early stages of development, so for us supporting such decision-making is the prevalent use of DDP. For example, in the study of an electronics packaging technique we identified 58 distinct Mitigations (Risk reducing actions in the form of preventions, detections, and alleviations) that could be performed to advance that technology from its current status as a research prototype to an engineering model. Selecting among them (given that it was unlikely that we could afford all 58 of them) was a necessity, and DDP supported this selection.

Clarify the purpose of risk reducing actions. For example, system tests may

uncover a wide variety of defects; however, if they are the primary means of validating system requirements, then the crafting and execution of those system tests should be approached with this in mind.

Support risk assessment. The direct assessment of the remaining risk can, on occasion, be difficult. We encounter this difficulty when challenged to assess technologies that themselves are novel. For example, the focus of one study was the potential use of MEMS (Micro Electrical Mechanical Systems) approaches to yield an exceptionally small, lightweight and low-power sensor that would operate in space environments. The combination of novel aspects of this application render direct risk assessment challenging.

We also encounter this difficulty when the technologies are relatively standard, but the applications are novel (for example, the aforementioned study of electronics for use on the Martian surface environment). In cases such as these there is insufficient past experience to assess risks directly. Instead, the DDP approach encourages the reconsideration of just how effective the various risk-reducing measures will be in these unusual circumstances.

Another setting where such concerns arise is in the assessment of software. Fenton et al motivate the need to reason over the various defect prevention, detection and correction steps taken during software development in order to assess the quality of the software that results [Fenton & Neil, 1999]. Their approach is to make use of Bayesian Belief Networks to combine knowledge of causal structure with evidence (expert judgments and/or historical records) of the error rates and efficacy of testing, etc. [Fenton et al, 2003].

5 Mitigations

In this section we consider DDP's representation of Mitigations – how it differs from conventional practices, and what are the comparative benefits that accrue from the DDP approach.

5.1 Conventional Practice

Conventional risk assessment processes vary considerably in the extent to which they explicitly represent Mitigations (actions that prevent or reduce risk). Probabilistic Risk Assessment approaches when used to assess the risk in a design normally do not explicitly represent any Mitigations, rather, they focus on assessing the risk remaining in that design. When risk assessment (e.g., using FMECAs) is performed in the earlier stages of development, it is common practice to identify the outstanding risks, and for the more significant risks, identify and recommend actions to take to reduce them. The effects of such actions are represented by describing the change to the risk that will occur when they are applied. Sometimes several such actions are considered together (i.e., only the net risk reduction of all those actions in combination is represented). Sometimes the sequential reduction of risk is represented, using a “ladder” diagram to show the risk being successively reduced as the actions are applied in sequence.

5.2 DDP's Treatment of Mitigations

In DDP, each Mitigation is assessed against each Risk, to capture the extent to which that Mitigation would reduce the Risk were it to be applied.

The benefits of DDP's treatment of Mitigations are as follows:

Support decision making to help choose which (out of potentially many) mitigations to employ. Discussion of this benefit began

earlier, in the section on Risks. This kind of decision making relies on being able to calculate the risk-reducing effects of arbitrary selections from among the Mitigations. Conventional risk assessment practices do not necessarily provide this information. For example, suppose we are told that the likelihood of a Risk is 0.6, and that the combination of two recommended actions will reduce this to 0.1; what would the risk-reducing effect be of just one of those actions? If we are given more information, say that the first of those recommended actions will reduce the likelihood from 0.6 to

0.2, and the second will further reduce this likelihood to 0.1, then we do know the effect of the first action alone, but we do not necessarily know the effect of the second action alone. While we might *assume* that since it halved the likelihood when applied after the first action, it would still halve the likelihood if applied on its own, i.e., reduce it from 0.6 to 0.3. DDP addresses such assumptions from the start: its treatment of Mitigations calls for them to be assessed individually, and there is an *explicit* assumption of how multiple Mitigations' effects on the same risk combine (notably that they act like filters in series – if one of them divides a Risk's likelihood by a factor of F1, and another by a factor of F2, then in combination they act to divide the Risk's likelihood by a factor of $F1 * F2$). When we have good reason to believe that Mitigations do *not* combine in this manner, we have some workarounds to accommodate this knowledge. For more discussion of DDP's calculations in this regard, see [Feather&Cornford 2003].

DDP's capabilities include a heuristic




Figure 3. DDP-generated cost/benefit trade space

search component to find near-optimal Mitigation selections [Cornford et al, 2003]. For example, suppose there is a known, limited budget for risk mitigation – optimization in this situation means find the selection of Mitigations that maximize attainment of Objectives while costing no more than that limit. An alternate approach is to set a minimum acceptable threshold of Objective attainment, and optimize to find the least-cost set of Mitigations that attain at least that much of the Objectives.

Reveal the overall trade space for cost vs. benefit.

An example of DDP's calculation and visualization of the overall trade space for one of our technology studies is shown in Fig. 3. The black “cloud” consists of hundreds of thousands of points, each representing a selection of Mitigations. The location of each point with respect to the horizontal axis (cost) is determined by the cost of that set of Mitigations (as calculated by DDP), and location with respect to the vertical axis (benefit) by the benefit, i.e., attainment of Objectives (again, as calculated by DDP). The

sum total cost of all mitigations (approximately \$4,750,000) determines the rightmost value of the x-axis, and the sum total value of all objectives (approximately 3,600) determines the topmost value of the y-axis. The upper-left frontier of the cloud is thus the “optimal” boundary, also referred to as the “Pareto front” [Sen&Yang, 1998]. The search that generates these points is by design biased to concentrate on that boundary, so there are likely many points in the interior region not shown on this figure. The key insight we get from this figure is the location of the “sweet spot” for funding the development – in this study, this is centered slightly to the right of the \$1,000,000 mark. If the development is funded at significantly less than this amount, then the benefit attainment drops dramatically. Conversely, if the development is funded at significantly more than this amount, then the gain in benefit attainment is very minor. These insights can be used to guide the decision on level of funding (as was the case in this study), or, if sufficient funding is simply not available, to motivate the consideration of “descope” options (recall the discussion in the Risks section).

Capture the cost and benefit distinctions between Mitigations that prevent, detect or alleviate Risks.

In DDP Mitigations are subdivided into three categories: preventions, detections and alleviations. Preventions and detections decrease the *likelihood* of Risks occurring, while alleviations decrease the *severity* of Risks should they occur. The distinction between preventions and detections is that only the latter imply the need to repair a (detected) Risk, which will incur some additional cost (the cost of repair). DDP’s calculations take these factors into consideration.

For example, consider the planning of a software development effort:

- Adoption of a coding standard is a Risk *prevention* – it decreases the likelihood of certain kinds of Risks associated with confusion among multiple

developers. The cost associated with adopting a coding standard is the creation of that standard in the first place, and the training of the developers to make them aware of it.

- Conducting a peer-review of software code is a Risk *detection* – code problems (in DDP terms, “Risks”) may be discovered by the review, and then corrected. The benefit will be a reduction in the likelihood of problems remaining in the code. The cost will be the sum of the cost of performing the review, and the cost of correcting problems discovered by the review. Note that a problem’s repair cost depends on the nature of the problem, and on when it is discovered.

DDP’s calculations capture the cost-benefit reasoning that shows the net value of early-phase prevention and detection of Risks which, if uncovered in later phases (e.g., testing) would incur much greater cost. Software development cost considerations such as these are discussed, for example, in [Kaner, 1996]. Of course, hardware is prone to similar cost considerations.

Capture the effect of mitigations that increase risks. Mitigations may reduce some Risks, but increase others. For example, a vibration test is a “detection” kind of Mitigation (as discussed above) that can be used to detect the presence of defects, but may itself induce defects. In our DDP applications we have encountered this phenomenon in several guises:

- In our electronics packaging study, application of a coating to protect against damage, while decreasing the likelihood and/or severity of such damage, was judged to have the disadvantage of making rework of that circuit much harder – essentially it *aggravates* (leaves the likelihood unchanged, but increases the consequence of) the rework Risk.
- Some kinds of circuit components, while advantageous in certain respects, have the disadvantage of elevating the

likelihood of some problems – essentially they *induce* (increase the likelihood of) Risks.

- In a study that involved a choice among alternative power sources, those design alternatives were represented as an extreme form of Risk inducers. Some power source was needed – the only way to mitigate the “lack of power” Risk was through the choice of (at least one) of them! However, each alternative was connected by risk-inducing effect links to the Risks associated with that kind of power source. Those Risks had been assigned an a-priori likelihood of *zero*, meaning that they would come into play if and only if they were induced.

6 Human-Computer Interface

In this section we consider the human computer interface appropriate to DDP – how and why it differs from conventional practices, and what are the comparative benefits that accrue from the DDP approach.

6.1 Conventional Practice

Early-lifecycle risk methods differ significantly from late-lifecycle risk methods in the quantity and detail of risk information they involve.

- When Probabilistic Risk Assessment methods are applied to assess a detailed design, the fault trees and event sequence diagrams can be intricate and voluminous. Software tools that support PRA provide an interface appropriate to the construction and viewing of these (e.g., with continuation symbols to indicate that a single node on one page is expanded further on another). The results of PRA methods include measures of the overall system reliability at various confidence levels, “cut sets” (non-redundant combinations of events that together cause system failure) comprising lists of events and their overall likelihoods, and various sensitivity analyses (e.g., by how much

system reliability would increase if a particular event was rendered impossible).

- Early-phase risk methods (e.g., FMECAs) tend to deal with relatively modest numbers of risks. At this stage simple spreadsheets may suffice to hold the information. As risk management continues through the development it becomes important to keep track of risk status, planned dates for applying mitigations, etc. For these purposes database support may be appropriate to maintain the information. A commonly used interface to show overall risk status is to group the risks into a 2-dimensional chart – see Fig. 4, where the number in a cells indicates the number of risks that fall into that cell (a textual list of risks usually accompanies the chart).




Figure 4. Conventional risk grid chart.

6.2 DDP Interfaces

DDP offers a variety of visualizations to convey information about its quantitatively linked Objectives, Risks and Mitigations. Some of these visualizations are variants of those found in more traditional risk methods, while some are specific to DDP.

Overall connectivity among Objectives, Risks and Mitigations is portrayed in the topology visualization shown earlier in Fig. 2. From this view it is easy to see an “unlinked” item (e.g., an Objective that does not connect to any Risks). For some datasets, this view is also adept at revealing unbalanced portfolios (e.g., Risks that connect to a large number of




Figure 5. Topology visualization revealing an unbalanced portfolio of mitigations

Mitigations but connect to few, or no, Objectives. Such an example is seen in Fig. 5, taken from [Feather et al, 2003].

The *quantitative* connectivity information among Objectives, Risks and Mitigations is portrayed through (large) matrices, or, since these are generally sparse, through a more compact equivalent in which only the non-zero values are indicated.

The status of Risks is portrayed through bar charts, and through DDP's risk region chart. Fig. 6 shows an example of the latter. It differs from the conventional risk grid chart (Fig. 4) in two important respects:

- individual Risks are shown positioned with respect to the axes, rather than grouped into grid cells;
- the boundary between high, medium and low Risks (seen in the chart as the




Figure 6. DDP risk region chart

boundaries between the red, yellow and green regions) are lines of constant risk ("isorisk" contours!) – these appear as straight lines on the DDP chart because the axes are *log* scale. The stair-step boundary between regions of the conventional risk grid chart only approximates this. Indeed, it is possible in the conventional risk grid chart for a risk located in a high region to actually be a *lower* than one in a medium region. See Fig. 7 for a pathological example: this shows a demarcation between high and medium regions at the 0.5 lines; a risk at the (0.45, 0.95) point falls in the upper right corner of the yellow (medium) region, and has the computed risk value of 0.4275, while a risk at the (0.55, 0.55) point




Figure 7. A pathological example

falls in the lower left corner of the red (high) region, and has the computed risk value of 0.3025 – so although the first risk is actually higher than the second, it falls into a lower region than does the second.

DDP is able to generate both forms of risk charts from its risk data, and can even overlay the grid chart demarcations on top of the log-scale risk region chart.

Comparison of risk between different selections of mitigations can be seen via appropriate use of bar charts. Fig. 8 shows an example where each bar represents a different Risk, using red to show risk in both selections, black to show increase in risk, and yellow decrease in risk, when going from the first selection to the second. Here the bars have been sorted into descending order of the risk in the second selection.




Figure 8. Bar chart for risk comparison

DDP also uses bar charts to show status of the significance and attainment of individual Objectives, and to show status of the benefit the individual Mitigations can offer.

Overall cost/benefit trade space is the visualized thought the kind of chart shown earlier in Fig. 3. Individual cost and benefit values for a particular selection of Mitigations are shown in a small set of gauges, allowing the user to see on the fly the net effects of changing Mitigation selections.

7 Results

DDP has been applied to assess over a dozen advanced technologies intended for spacecraft application. These technologies have included hardware and software combinations such as an imaging technology, a miniaturized gyroscope, a compact data storage device, adaptation of a visual

programming language to spacecraft control, optical sensor technology, circuit board fabrication, a micro electrical mechanical system for determining the direction to the sun, and electronics that will be resistant to the extreme temperature swings of the Martian planetary environment.

DDP has also been used to assess and plan the risk mitigation strategies for an entire space experiment, and is in current use as the overall risk management approach for an ongoing flight mission that is in its early phases of development and planning.

These applications have generally led to improvements in the plans for how to continue the development of the studied technologies. The nature of these results varies from case to case. For example, the study of the micro electrical mechanical system for determining the direction to the sun led to:

- A clearer definition of the work needed to mature the technology towards spacecraft use;
- Identification of a commercial opportunity based on the exceptional performance offered by the new technology;
- Improvements over the initial estimates of cost to complete the development, by considering the key tasks that are spacecraft project specific in nature;
- Determination of a design that minimizes risk specifically for the intended spacecraft application.

In some instances, the outcome of the studies has been dramatically improved designs – these improvements have been in areas of cost and other spacecraft-critical resources (e.g., mass, electrical power). The cost of a DDP study of a technology is typically in the range of \$10,000 to \$30,000 (paying for the time of the discipline area experts involved in the study). The improved designs reflect savings that far exceed these costs.

8 Conclusions

We have indicated where and how the DDP approach differs from most other forms of risk assessment. The utility of the DDP approach has served us well in applications to individual technologies intended for spacecraft applications, and to ongoing risk assessment and risk management for the early phases of an entire mission. This is not to say that DDP is the perfect solution. DDP has drawbacks in two main areas:

Significant effort is required to gather the information that DDP works with. This is especially noticeable in comparison with other early-lifecycle risk assessment practices, which do not require the identification of Objectives, or the detailed assessment of individual Mitigations' effects at reducing risk. Thus the application of DDP comes at a price, namely the time and effort it takes to gather the information required.

Shortcomings remain in the DDP model. DDP lacks, for example, the ability to deal with distributions and uncertainty, with complex models of utility (e.g., an accuracy objective whose attainment is calculated as, say, the root mean square of the accuracy attainment of its components), and with the fault-tree gates and event-sequence diagram constructs that PRA techniques employ to capture the nuances of detailed designs. We have begun to address some of these shortcomings (e.g., we have incorporated logical fault trees with "and" and "or" gates into the DDP model of Risks). We have proposed to use DDP in an iterative cooperation with PRA, to use DDP home in on the areas where the more in-depth application of PRA is warranted [Cornford et al, 2003]. We are also beginning to investigate a direct connection between DDP and the Galileo PRA tool [Sullivan et al, 1999].

9 References

[Cornford, 1998] S.L. Cornford. "Managing Risk as a Resource using the

Defect Detection and Prevention process" *4th International Conference on Probabilistic Safety Assessment and Management*, 13-18 September 1998, New York City, NY, International Association for Probabilistic Safety Assessment and Management.

[Cornford et al, 2001] S.L. Cornford, M.S. Feather & K.A. Hicks. "DDP – A tool for life-cycle risk management", *IEEE Aerospace Conference*, Big Sky, Montana, Mar 2001, pp. 441-451.

[Cornford et al, 2003] S.L. Cornford, M.S. Feather, J. Dunphy, J. Salcedo & T. Menzies, 2002, "Optimizing the Design of end-to-end Spacecraft Systems using Risk as a Currency", *IEEE Aerospace Conference*, Big Sky, Montana, Mar 2003, pp. 7.3361 – 7.3368.

[Cornford et al, 2003b] S.L. Cornford, T. Paulos, L. Meshkat & M.S. Feather. "Towards More Accurate Life Cycle Risk Management Through Integration of DDP and PRA", *IEEE Aerospace Conference*, Big Sky, Montana, Mar 2003, pp. 2.1106 – 2.1200.

[Feather et al, 2000] M.S. Feather, S.L. Cornford & M. Gibbel. "Scalable Mechanisms for Goals Interaction Management", *4th IEEE International Conference on Requirements Engineering*, Schaumburg, Illinois, 19-23 Jun 2000, IEEE Computer Society, pp 119-129

[Feather et al, 2002] M.S. Feather, S.L. Cornford & K.A. Hicks. "Descoping". *27th NASA IEEE Software Engineering Workshop*, Greenbelt Maryland, Dec 2002.

[Feather&Cornford, 2003] M.S. Feather. & S.L. Cornford. "Quantitative risk-based requirements reasoning", *Requirements Engineering* (Springer), Vol 8 #4, pp 248-265, 2003; published online 25 February 2003, DOI 10.1007/s00766-002-0160-y.

[Feather et al, 2003] M.S. Feather, T. Menzies & J.R. Connelly. "Matching Software Practitioner Needs to Researcher Activities", *2003 Asia-Pacific Software Engineering Conference (APSEC 2003)*; Chiangmai, Thailand, Dec. 2003, pp. 6-16.

[Fenton & Neil, 1999] N. Fenton & M. Neil. "A Critique of Software Defect

Prediction Research”, *IEEE Transactions on Software Engineering* 25(5), 1999.

[Fenton et al, 2003] N. Fenton, P. Krause & M. Neil “A Probabilistic Model for Software Defect Prediction”, To appear in *IEEE Transactions on Software Engineering* – contact Fenton at: norman@dcs.qmw.ac.uk

[Kaner, 1996] C. Kaner. “Quality Cost Analysis: Benefits and Risks”, *Software QA* Vol 3, #1, p. 23, 1996.

[Sen&Yang, 1998] P. Sen & J-B. Yang. *Multiple Criteria Decision Support in Engineering Design*, Springer-Verlag, 1998.

[Sullivan et al, 1999] K.J. Sullivan, J.B. Dugan & D. Coppit. “The Galileo fault tree analysis tool”, *29th Annual International Symposium on Fault-Tolerant Computing*, Madison, Wisconsin, June 1999, pp. 232-235.